Search Results

Search Results for: **[(pareto and task and concurrency)]**
Found **17** of **126,861 searched.**

Search within Results

> Advanced Search

> Search Help/Tips

Sort by: **Title** **Publication** **Publication Date** **Score** Binder

Results 1 - 17 of 17 **short listing**

1 Invited talk: Managing dynamic concurrent tasks in embedded real-time multimedia systems 94%
Peng Yang , Paul Marchal , Chun Wong , Stefaan Himpe , Francky Catthoor , Patrick David , Johan Vounckx , Rudy Lauwereins
Proceedings of the 15th international symposium on System Synthesis October 2002
This paper addresses the problem of mapping an application, which is highly dynamic in the future, onto aheterogeneous multiprocessor platform in an energy efficient way. A two-phase scheduling method is used for that purpose. By exploring the Pareto curves and scenarios generated at design time, the run-time scheduler can easily find a good scheduling at a very low overhead, satisfying the system constraints and minimizing the energy consumption. A real-life example from a 3D quality of service ...

2 Session 5A: Embedded tutorial: embedded software and systems: Optimisation problems for dynamic concurrent task-based systems 89%
D. Verkest , P. Yang , C. Wong , P. Marchal
Proceedings of the 2001 IEEE/ACM international conference on Computer-aided design November 2001

3 Task concurrency management methodology to schedule the MPEG4 IM1 player on a highly parallel processor platform 88%
Chun Wong , Paul Marchal , Peng Yang
Proceedings of the ninth international symposium on Hardware/software codesign April 2001

This paper addresses the concurrent task management of complex multi-media systems, like the MPEG4 IM1 player, with emphasis on how to derive energy-cost *vs* time-budget curves through task scheduling on a multi-processor platform. Starting from the original "standard" specification, we extract the concurrency originally hidden by implementation decisions in a "grey-box" model. Then we have applied two high-

level transformations on this model to improve the task- ...

4 Effective distribution of object-oriented applications 82%
Sandeep Purao , Hemant Jain , Derek Narareth
Communications of the ACM August 1998
Volume 41 Issue 8

5 Persistent storage for a workflow tool implemented in Smalltalk 82%
Bob Beck , Steve Hartley
ACM SIGPLAN Notices , Proceedings of the ninth annual conference on Object-oriented programming systems, language, and applications October 1994
Volume 29 Issue 10
This paper describes a new workflow model and its implementation in Smalltalk. The paper also details problems with using a RDBMS as the persistent store for the workflow tool and the subsequent experiences in using an ODBMS for this purpose. The final solution was a coexistence approach, using the RDBMS for legacy corporate data and the ODBMS for the process description and workflow status data.

6 Minnowbrook APL workshop 80%
R H Pesch , E E McDonnell , K E Iverson , B Bernecky , D B Allen
ACM SIGAPL APL Quote Quad March 1986
Volume 16 Issue 3

7 RS-FDRA: a register sensitive software pipelining algorithm for embedded VLIW processors 80%
Cagdas Akturan , Margarida F. Jacome
Proceedings of the ninth international symposium on Hardware/software codesign April 2001

The paper proposes a novel software-pipelining algorithm, *Register Sensitive Force Directed Retiming Algorithm (RS-FDRA)*, suitable for optimizing compilers targeting embedded VLIW processors. The key difference between RS-FDRA and previous approaches is that our algorithm can handle *code size constraints* along with latency and resource constraints. This capability enables the exploration of pareto "optimal" points with respect to *code size* and *performance*

8 Low-power systems on chips (SOCs) 80%
C. Piguet , M. Renaudin , T. Omnés
Proceedings of the conference on Design, automation and test in Europe March 2001

9 Predicting parallel applications performance on non-dedicated cluster platforms 80%
Cosimo Anglano
Proceedings of the 12th international conference on Supercomputing July 1998

10 Design space exploration algorithm for heterogeneous multi-processor embedded system design 80%
Ireneusz Karkowski , Henk Corporaal
Proceedings of the 35th annual conference on Design automation conference May

1998

11 Measurement and analysis of sequential design processes 80%

E. W. Johnson , J. B. Brockman

ACM Transactions on Design Automation of Electronic Systems (TODAES) January 1998

Volume 3 Issue 1

As design processes continue to increase in complexity it is important to base process-improvement decisions on quantitative analysis. We describe the development of an analytical approach for evaluating sequential design-process completion time and for determining the sensitivities of design time with respect to individual task durations and transition probabilities. Techniques are also detailed for collecting process metadata and calibrating a design proecess model. Example applications i ...

12 SOC test architecture design for efficient utilization of test bandwidth 77%

Sandeep Kumar Goel , Erik Jan Marinissen

ACM Transactions on Design Automation of Electronic Systems (TODAES) October 2003

Volume 8 Issue 4

This article deals with the design of on-chip architectures for testing large system chips (SOCs) for manufacturing defects in a modular fashion. These architectures consist of wrappers and test access mechanisms (TAMs). For an SOC with specified parameters of modules and their tests, we design an architecture that minimizes the required tester vector memory depth and test application time. In this article, we formulate the test architecture design problems for both modules with fixed- and flexi ...

13 Cluster resource management: Integrated resource management for cluster-based Internet services 77%

Kai Shen , Hong Tang , Tao Yang , Lingkun Chu

ACM SIGOPS Operating Systems Review December 2002

Volume 36 Issue SI

Client request rates for Internet services tend to be bursty and thus it is important to maintain efficient resource utilization under a wide range of load conditions. Network service clients typically seek services interactively and maintaining reasonable response time is often imperative for such services. In addition, providing differentiated service qualities and resource allocation to multiple service classes can also be desirable at times. This paper presents an integrated resource managem ...

14 Real-time FRP 77%

Zhanyong Wan , Walid Taha , Paul Hudak

ACM SIGPLAN Notices , Proceedings of the sixth ACM SIGPLAN international conference on Functional programming October 2001

Volume 36 Issue 10

Functional reactive programming (FRP) is a declarative programming paradigm where the basic notions are continuous, time-varying behaviors and discrete, event-based reactivity. FRP has been used successfully in many reactive programming domains such as animation, robotics, and graphical user interfaces. The success of FRP in these domains encourages us to consider its use in real-time applications, where it is crucial that the cost of running a program be bounded and known before run-time. But p ...

15 Task concurrency management methodology summary 77%

C. Wong , P. Marchal , P. Yang , F. Catthoor , H. de Man , A. Prayati , N. Cossement , R. Lauwereins , D. Verkest
Proceedings of the conference on Design, automation and test in Europe March 2001

16 Process migration 77%

ACM Computing Surveys (CSUR) September 2000
Volume 32 Issue 3

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

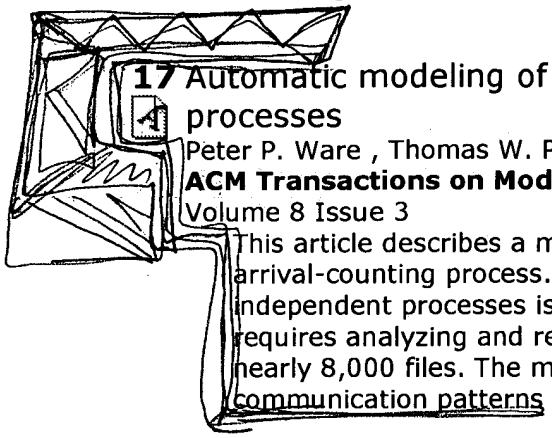

17 Automatic modeling of file system workloads using two-level arrival processes 77%

Peter P. Ware , Thomas W. Page , Barry L. Nelson
ACM Transactions on Modeling and Computer Simulation (TOMACS) July 1998
Volume 8 Issue 3

This article describes a method for analyzing, modeling, and simulating a two-level arrival-counting process. This method is particlarly appropriate when the number of independent processes is large, as is the case in our motivating application which requires analyzing and representing computer file system trace data for activity on nearly 8,000 files. The method is also applicable to network trace data characterizing communication patterns between pairs of computers. We apply cluster analy ...

Results 1 - 17 of 17 **short listing**

Search Results

Search Results for: **[(task concurrency management)]**
Found **9** of **126,861 searched.**

Search within Results

Go > Advanced Search

> Search Help/Tips

Sort by: **Title** **Publication** **Publication Date** **Score** Binder

Results 1 - 9 of 9 **short listing**

1 Session 5A: Embedded tutorial: embedded software and systems: Optimisation problems for dynamic concurrent task-based systems 88%
D. Verkest , P. Yang , C. Wong , P. Marchal
Proceedings of the 2001 IEEE/ACM international conference on Computer-aided design November 2001

2 Efficient system exploration and synthesis of applications with dynamic data storage and intensive data transfer 88%
Julio Leao da Silva , Chantal Ykman-Couvreur , Miguel Miranda , Kris Croes , Sven Wuytack , Gjalt de Jong , Francky Catthoor , Diederik Verkest , Paul Six , Hugo De Man
Proceedings of the 35th annual conference on Design automation conference May 1998

Matisse is a design flow intended for developing embedded systems characterize dby tight inter action b etwe encontrol and data-flow behavior, intensive data storage and tr ansfer, dynamic creation of data, and stringent real-time requirements. Matisse bridges the gap from a system specification, using a cocurr ent obje ct-oriented language, to an optimize d embedded single-chip HW/SW implementation. Matisse supp orts stepwise system-level exploration and refinement, memory architec ...

3 Task concurrency management methodology to schedule the MPEG4 IM1 player on a highly parallel processor platform 87%
Chun Wong , Paul Marchal , Peng Yang
Proceedings of the ninth international symposium on Hardware/software codesign April 2001

This paper addresses the concurrent task management of complex multi-media systems, like the MPEG4 IM1 player, with emphasis on how to derive energy-cost *vs* time-budget curves through task scheduling on a multi-processor platform. Starting from the original "standard" specification, we extract the concurrency originally hidden by implementation decisions in a "grey-box" model. Then we have applied two high-

level transformations on this model to improve the task- ...

4 Task concurrency management methodology summary 87%

C. Wong , P. Marchal , P. Yang , F. Catthoor , H. de Man , A. Prayati , N. Cossement , R. Lauwereins , D. Verkest

Proceedings of the conference on Design, automation and test in Europe March 2001

5 Invited talk: Managing dynamic concurrent tasks in embedded real-time multimedia systems 82%

Peng Yang , Paul Marchal , Chun Wong , Stefaan Himpe , Francky Catthoor , Patrick David , Johan Vounckx , Rudy Lauwereins

Proceedings of the 15th international symposium on System Synthesis October 2002

This paper addresses the problem of mapping an application, which is highly dynamic in the future, onto aheterogeneous multiprocessor platform in an energy efficient way. A two-phase scheduling method is used for that purpose. By exploring the Pareto curves and scenarios generated at design time, the run-time scheduler can easily find a good scheduling at a very low overhead, satisfying the system constraints and minimizing the energy consumption. A real-life example from a 3D quality of service ...

6 Practical experiences: Multiprocessor mapping of process networks: a JPEG decoding case study 82%

E. A. de Kock

Proceedings of the 15th international symposium on System Synthesis October 2002

We present a system-level design and programming method for embedded multiprocessor systems. The aim of the method is to improve the design time and design quality by providing a structured approach for implementing process networks. We use process networks as re-usable and architecture-independent functional specifications. The method facilitates the cost-driven and constraint-driven source code transformation of process networks into architecture-specific implementations in the form of communi ...

7 Task scheduling and real-time: System-level power-performance trade-offs in task scheduling for dynamically reconfigurable architectures 80%

Juanjo Noguera , Rosa M. Badia

Proceedings of the international conference on Compilers, architectures and synthesis for embedded systems October 2003

Dynamic scheduling for System-on-Chip (SoC) platforms has become an important field of research due to the emerging range of applications with dynamic behavior (e.g. MPEG-4). Dynamically reconfigurable architectures are an interesting solution for this type of applications.However, dynamic scheduling for run-time reconfigurable architectures with power-performance trade-offs has not been addressed in previous research efforts. In this paper, we address this open issue using a system-level approa ...

8 System design methods: scheduling advances: Dynamic run-time HW/SW scheduling techniques for reconfigurable architectures 80%

Juanjo Noguera , Rosa M. Badia

Proceedings of the tenth international symposium on Hardware/software codesign May 2002

Dynamic run-time scheduling in System-on-Chip platforms has become recently an active area of research because of the performance and power requirements of new applications. Moreover, dynamically reconfigurable logic (DRL) architectures are an exciting alternative for embedded systems design. However, all previous approaches to DRL multi-context scheduling and HW/SW scheduling for DRL architectures are based on static scheduling techniques. In this paper, we address this problem and present: (1) ...

9 Proposal for unified system design meta flow in task-level and instruction-level design technology research for multi-media applications 80%

Francky Catthoor , Diederik Verkest , Erik Brockmeyer

Proceedings of the 11th international symposium on System synthesis December 1998

Results 1 - 9 of 9 short listing